

FIWARE FOR DIGITAL TWINS

POSITION PAPER / VERSION 1.0 / JUNE 2021

VISIT...

LANZAROTE
Caliente.COM

Publisher

FIWARE Foundation e.V.
Franklinstraße 13A
10587 Berlin
Germany

Copyright

FIWARE Foundation e.V., June 2021



Editor

Juanjo Hierro, FIWARE Foundation

Contributors

This paper has been revised under the umbrella of the FIWARE Digital Twin Roadmap WG activities in collaboration with ETSI CIM ISG. The FIWARE Foundation wishes to thank the following members of the WG who have participated in revision of the document.

Alberto Abella, FIWARE Foundation
Álvaro Alonso, UPM
Martin Bauer, NEC
Javier Conde, UPM
Lindsay Frost, NEC
Franck Le Gall, EGM
Benoit Orihuela, EGM
Gilles Privat, Orange
Joaquín Salvachúa, UPM
Giuseppe Tropea, CNIT
Ken Zangelin, FIWARE Foundation

Disclaimer

In accordance with our Guidelines concerning the use of technology, strategic and market information in positioning and advertising, please be aware of the following: Position Papers appearing on the FIWARE Foundation site or in other digital or printed materials are actually received via text, audio or video submission. They are individual experiences and common strategies, reflecting studies and real life experiences of those who will use or have used our technology and/or services in some way or another. We do not claim that they are typical results that customers will generally use and achieve. Some FIWARE Position Papers have been shortened.

Table of contents

Overview.....	4
1. FIWARE and Digital Twins.....	4
2. Vertical Smart Solutions powered by FIWARE	8
3. Smart Organizations powered by FIWARE.....	10
4. Data Spaces powered by FIWARE.....	12
5. The Smart Data Models initiative.....	15
Conclusions	17

Overview

This paper describes how smart applications from multiple domains can participate in the creation of smart solutions based on the Digital Twin paradigm using FIWARE software building blocks. Central in the vision, the NGSI-LD API is proposed as an open standard

API for getting access to digital twin data and use of standard data models is promoted to ensure portability and replicability of solutions. Smart solutions powered by FIWARE are architected around management of a digital twin data representation of the real world sharing their data.

1 FIWARE and Digital Twins

[FIWARE](#) was created with the ultimate goal of creating an open sustainable ecosystem around public, royalty-free and implementation-driven software platform standards easing the development of smart solutions and supporting organizations in their transition into smart organizations. From a technical perspective, FIWARE brings a curated framework of Open Source software components which can be assembled together and combined with other third-party platform components to build platforms easing the development of smart solutions and smart organizations in multiple application domains: cities, manufacturing, utilities, agrifood, etc. Since creation of the FIWARE Foundation in late 2016, a vibrant FIWARE Community has been formed with a true worldwide dimension, comprising [more than 90 member organizations](#), including large corporations, SMEs, technology centres and universities, and hundreds of in-

dividual members. Parallel to this growth the number of organizations adopting FIWARE has not stopped increasing.

Any software architecture “powered by FIWARE” gravitates around management of a Digital Twin data representation of the real world. This Digital Twin data representation is built based on information gathered from many different sources, including sensors, cameras, information systems, social networks, end users through mobile devices, etc. It is constantly maintained and accessible in near real-time (“right-time” is the term also often used, reflecting that the interval between the instants of time at which some data is gathered and made accessible is enough short to allow a proper reaction). Applications constantly process and analyze this data (not only current values but also history generated




Figure 1 - Current member organizations in the FIWARE Community

over time) in order to automate certain tasks or bring support to smart decisions by end users. The collection of all Digital Twins modelling the real world that is managed is also referred to as **Context** and the data associated with attributes of Digital Twins is also referred to as **context information**.

In FIWARE, a Digital Twin is an entity which digitally represents a real-world physical asset (e.g. a bus in a city, a milling machine in a factory) or a concept (e.g., a weather forecast, a product order). Each Digital Twin:

- is universally identified with a URI (Universal Resource Identifier),

- belongs to a well-known type (e.g., the Bus type, of the Room type) also universally identified by a URI, and
- is characterized by several attributes which in turn are classified as:
 - properties holding data (e.g., the “current speed” of a Bus, or “max temperature” in a Room) or
 - relationships, each holding a URI identifying another Digital Twin entity the given entity is linked to (e.g., the concrete Building where a concrete Room is located).

Attributes of a Digital Twin may vary ranging from attributes that are quite static (e.g., the “li-




Figure 2 - Levels of Integration supported following a Digital Twin approach

cense plate” of a Bus), to attributes that change very dynamically (e.g., the “speed” or “number of passengers” in a Bus) to attributes which still change but not that often (e.g., the “driver” in a Bus which may change twice a day). Very important, the attributes of a Digital Twin are not only limited to observable data but also inferred data. Thus, for example, the Digital Twin of a Street may not only have an attribute “current traffic”, which may be automatically measured through sensors or cameras, but an attribute “forecasted traffic in 30 minutes” which might be calculated based on AI algorithms that keep the value of this attribute updated based on current traffic data, other relevant data impacting traffic (e.g., current weather observed and forecasted, schedule of events nearby, etc) and historical information about traffic in the given street. Therefore, the Digital Twin data representation of the world that is managed in a “powered by FIWARE” architecture is expected to contain all the information needed by smart

applications, not only measurable data but also other augmented insights and knowledge acquired over time.

A Digital Twin approach provides the basis for data integration at different levels, as illustrated in *Figure 2*:

- Within a **vertical Smart Solution**, solving how main Building Blocks within the architecture can be integrated;
- Within a **Smart Organization**, bringing support to the integration of the different systems within a smart organization following a system of systems approach;
- Within a **Smart Data Space**, establishing the basic “common lingua” that systems linked to the different organizations speak and understand.

Two critical elements need to be standardized in order to support an effective data integration at these three levels: the API to get access to Digital Twin data and the data models describ-

ing the attributes and semantics associated with the different types of Digital Twins being considered. The FIWARE Community has driven and continues to drive standardization at both fronts:

- The **NGSI API** provides a simple yet powerful RESTful API for getting access to context / Digital Twin data. NGSI API specifications have evolved over time driven by feedback from developers and implementation experiences. A first mature version of the API was the **NGSiv2 API**, which was defined by members of the FIWARE Community and is currently used in many systems in production within multiple sectors. Evolution of the API has taken place within the [ETSI ISG CIM](#) (Context Information Management Industry Specification Group), where members of the FIWARE Community and the FIWARE Foundation have led the definition of an evolved version of the API, known as the [NGSI-LD API](#), whose specifications were first published by ETSI in 2019 and continue to evolve. The NGSI-LD API is used as the data integration API and is implemented by the core component of any “powered by FIWARE” architecture: the so-called Context Broker component. Different alternative Open Source implementations of a Context Broker are available within the FIWARE Community, namely the Orion-LD, Scorpio and Stellio products.
- The **Smart Data Models initiative** ([website](#), [github](#)), launched by the FIWARE Foundation, provides a library of Data Models described in JSON/JSON-LD format which are compatible respectively with the NGSiv2/NGSI-LD APIs or would be useful for defining other RESTful interfaces for accessing Digital Twin data.

Data models published under the initiative are compatible with schema.org and comply with other existing de-facto sectoral standards when they exist. They solve one major issue developers are facing, that is the fact that a given data model specification may be mapped into JSON/JSON-LD in many different ways, all of them valid. Thanks to the Smart Data Models initiative, developers can rely on concrete mappings into JSON/JSON-LD, compatible with the NGSiv2/NGSI-LD APIs, that are made available within this library, avoiding interoperability problems derived from alternative mappings. Since its creation, more than 500 data models have been published and the number of organizations contributing data model descriptions is constantly growing. Relevant organizations like [TM Forum](#), [OASC](#) or [IUDX](#) are joining forces with the FIWARE Foundation bringing support to an open governance model for the initiative, following best Open Source practices.

Building around the FIWARE Context Broker, which is the core mandatory component in any “powered by FIWARE” architecture, a rich set of complementary Open Source components are available listed as part of the [FIWARE Catalogue](#). These components can be classified in the following categories or chapters as illustrated in *Figure 3*:

- components easing development of interfaces with the Internet of Things, Robotic and third-party systems;
- components supporting context / Digital Twin data processing and monitoring or the connection with data processing engines (e.g., Apache Spark, Apache Flink), monitoring tools (e.g., Grafana) and analysis platforms (e.g., Apache Superset);




Figure 3 - FIWARE Catalogue chapters

- components covering aspects related to Identity and Access Management (IAM) as well as Publication and Monetization of Data (including data accessible via APIs like NG-SI-LD).

2 Vertical Smart Solutions powered by FIWARE

Figure 4 depicts the reference architecture of a vertical smart solution powered by FIWARE. The concrete example corresponds to a smart solution for picking and palletizing products from a warehouse using robots.

This reference architecture is structured in essentially three layers:

- A Context Broker component is at the core of the architecture, keeping a digital twin representation of the real world objects and




Figure 4 - Smart Solution for Picking and Palletizing products from a warehouse

concepts relevant to the specific problem tackled: AGV robots, palletizer robots, shelf sections where products are stored in the warehouse, automatic doors AGV robots have to pass, operators in the shopfloor, items of stored products, orders generated from the CRM system, etc.

- Southbound to the Context Broker, the NGSI IoT Agents, available as part of the FIWARE IDAS framework, are used for connections to robotic systems exporting the OPC-UA IoT protocol or to specific sensors or actuators, used for example to detect items in shelf sections or to be able to open the shop floor

doors. They perform the necessary conversions between IoT protocols and NGSI. In addition, System Adapters developed based on the IDAS Agent library cope with the connection to the CRM and the Warehouse Inventory Management system that the solution has to interface with. FIWARE components like FIROS/SOSS allow, on the other hand, to perform the adaptation to robotic systems based on ROS/ROS2. Last but not least, the FIWARE component Kurento is able to process the video streams of cameras deployed in the shop floor, which are helpful to detect potential obstacles or risky situations.

- Northbound to the Context Broker, a number of tools are targeted to support real-time big data processing of the streams of history data generated as context / digital twin information evolves over time. A combination of open source components from third party products (Apache Superset, Grafana) and FIWARE (e.g., Wirecloud) are shown in the picture targeted to support the creation of operational dashboards and advanced data maps for monitoring processes. A number of FIWARE Data Connectors (Cygnus, Draco, Cosmos, STH Comet, QuantumLeap) are available as part of FIWARE to facilitate transference of historic context / digital twin information to these tools.

Transversal to all these layers, a number of FIWARE components support Identity and Access Management (e.g., Keyrock, API Umbrella, AuthZ-Force). They control the flow of data across the different layers. With regards to the access to the Con-

text Broker, they enforce the policies establishing what users can update, query or subscribe to changes on context / digital twin data. Note that the flow of data is not only south to north in the picture. Northbound applications can perform updates on context data, which in turn will trigger changes in the devices, robots or systems that are connected southbound.

An important point to highlight is that FIWARE is not about taking it all or nothing. You are not forced to use all the complementary FIWARE components mentioned above but you are free to use other third party platform components as well to design the hybrid platform of your choice. Thus, for example, you may opt for using a concrete IoT platform instead of IDAS IoT Agents to interface with sensors and actuators as reflected in the picture. As long as it uses the FIWARE Context Broker technology to manage context information, your platform can be labeled as “powered by FIWARE” and solutions built on top as well.

3 Smart Organizations powered by FIWARE

FIWARE can help organizations to digitally transform themselves and become smart organizations, bringing support to the integration of the different systems within the organization following a system of systems approach.

Figure 5 depicts the reference architecture of a smart city powered by FIWARE, for example. Again, the Context Broker component is at the

core of the architecture, holding a digital twin representation of the real world objects and concepts and describing what is going on in the city: streets, waste bins and containers, waste trucks, buses, electric vehicle chargers, buildings, events, citizen claims, etc. The different vertical smart solutions deployed in the city (e.g., Air Quality Monitoring, Smart Traffic Management, Smart Parking, Smart Waste Manage-




Figure 5 - Smart City Reference Architecture

ment) are connected to the Context Broker contributing the information they manage, which is relevant for creating a holistic Context / Digital Twin representation of the whole city, thereby breaking the information silos. Some of these vertical smart solutions may be powered by FIWARE (e.g., Traffic Control and Waste Management systems in the figure) in which case their interface with the global city-level Context Broker does not require any adaptation. Others may not be powered by FIWARE but this doesn't represent a major problem because creation of NGSI system adapters which translate from whatever API those systems export to NGSI-LD has proven not to be difficult. Last but not least, the city may deploy sensor/camera infrastructures through which valuable data is extracted.

Exploiting the complete Context / Digital Twin representation of the city, the Smart City

Governance System (or City Operation Center) can be developed. Real-time BigData processing tools can be used relying on data coming from multiple sources, extracting more valuable insights for the support of decisions. Similarly, monitoring tools can leverage this holistic Context / Digital Twin representation of the city.

Completing the picture, FIWARE brings components which provide the means for tracing digital twin data transactions. This provides the basis for a number of important functions, from identification of the provenance of data to audit-proof logging of certain transactions. For those organizations with strong requirements on transparency and certification, FIWARE brings components (i.e., **Canis Major**) that ease recording of transaction logs into different **Distributed Ledgers / Blockchains**.

4 Data Spaces powered by FIWARE

A **Data Space** can be defined as a decentralized data ecosystem built around commonly agreed Building Blocks enabling an effective and trusted sharing of data among participants. From a technical perspective, a number of **technology Building Blocks** are required ensuring:

- **Data interoperability.** Data Spaces should provide a solid framework for an efficient exchange of data among participants, supporting full decoupling of data providers and consumers. This requires the adoption of a “common lingua” every participant uses, materialized in the adoption of common APIs for the data exchange, and the definition of common data models. Common mechanisms for traceability of data exchange transactions and data provenance, are also required.
- **Data sovereignty and trust.** Data Spaces should bring technical means for guaranteeing that participants in a Data Space can trust each other and exercise sovereignty over data they share. This requires the adoption of common standards for managing the identity of participants, the verification of their truthfulness and the enforcement of policies agreed upon data access and usage control.
- **Data value creation.** Data Spaces should provide support for the creation of multi-sided markets where participants can generate value out of sharing data (i.e., creating data value chains). This requires the adoption of common mechanisms enabling the definition of terms and conditions (including pricing) linked to data offerings, the publication and discovery of such offerings and the management of all the

necessary steps supporting the lifecycle of contracts that are established when a given participant acquires the rights to access and use data.

Besides the adoption of a common technology foundation, Data Spaces also require **governance**, that is the adoption of a number of business, operational and organizational agreements among participants. Business agreements, for example, specify what kind of terms and conditions can regulate the sharing of data between participants and the legal framework supporting contracts established through the Data Space. Operational agreements, on the other hand, regulate policies that have to be enforced during Data Space operation like, for example, compliance with GDPR (General Data Protection Regulation) or the 2nd Payment Services Directive (PSD2) in the finance sector. They may also comprise the definition of tools that operators of cloud infrastructures or global services supporting Data Spaces must implement, enabling auditing of certain processes or the adoption of cyber-security practices. Last but not least, organizational agreements establish the governance bodies (very much like ICANN for the Internet). They deal with the identification of concrete specifications that products implementing technology Building Blocks in a Data Space should comply with, as well as the business and operational agreements to be adopted. The complete taxonomy of Building Blocks required for creating Data Spaces is illustrated in *Figure 6*. This same taxonomy of building blocks is also described in the [Open DEI White Paper on Design Principles for Data Spaces](#).




Figure 6 - Building Blocks in a Data Space

Sharing of data within a given Data Space should not be limited to a single domain. This would severely limit the creation of new innovative services since individuals and organizations usually act in multiple domains at the same time and many opportunities will flourish when data generated within organizations operating in certain domain (management of traffic in cities, for example) is shared for its exploitation in processes relevant to other domains (continuing with the example, logistics). Therefore, technology Building Blocks for Data Spaces must be domain-agnostic. On the other hand, they should rely on open standards, allowing multiple infrastructure and global service providers to emerge and support Data Spaces, without getting locked by any particular provider. Given this, while making things work in living labs and pilots is relatively easy, the main challenge towards definition of successful Data Spaces is the decision of what concrete standards and design principles are adopted, since they have to be accepted by all participants.

The following sections elaborate on the different components FIWARE brings materializing the dif-

ferent technical Building Blocks required for creation of Data Spaces. The [i4Trust initiative](#) is creating a curated framework based on these components that will be available for creation of Data Spaces and pioneers experiments during the period 2021-2023.

Data providers joining data spaces must be able to publish data resources at well defined endpoints knowing that data consumers, a priori unknown to them, will know how to retrieve and consume data through those endpoints. Data consumers, on the other hand, must know how data available through endpoints they discover can be consumed. This is a key principle which was observed in the design of the world wide web: content providers publish web pages on web servers (endpoints) knowing that web browsers will be able to connect to them and retrieve web pages whose content they can render and display to end users. It means that all participants in data spaces should 'speak the same language', which translates into adopting domain-agnostic common APIs and security schemas for data exchange (the way of constructing sentences) to-




Figure 7 - Data Exchange in a Data Space “powered by FIWARE”

gether with data models represented in data formats compatible with those APIs (the vocabulary used in constructed sentences).

The **NGSI API** is domain-agnostic. Actually, many different systems have been developed using NGSI in domains such as Smart Cities, Smart Manufacturing, Smart Energy, Smart Water, Smart AgriFood, Smart Ports, or Smart Health, to mention a few. This facilitates data sharing because each system participating in a data space will be publishing data that simply enriches a Digital Twin data representation of the world that the rest of systems connecting to the data space will know how to access. Systems participating in the data space don’t know a priori what other systems may consume the data they publish (although they will be able to set up concrete terms and conditions for accessing/using data as we will explain in the next section).

Figure 7 illustrates how different systems participating in data spaces “powered by FIWARE” will exchange data. Context Broker servers are the endpoints through which systems connected to the data space publish digital twin data, very much like web servers publish html content on the world wide web. Those systems can in turn connect to Context Broker servers in order to obtain information they need. Note that data spaces powered by FIWARE enable near real-time (right-time) exchange of digital twin data which is fundamental in the design of innovative value chains demanding a very dynamic exchange of data among participants. Just think about scenarios like a city managing traffic lights in streets close to a given train station in order to facilitate that travelers arriving and taking a taxi can leave faster to their destinations. NGSI-LD brings very simple and therefore easy to use operations for creating, updating and consuming context / digital twin data but also more powerful operations like sophisticated que-

ries, including geo-queries, or the subscription to get notified on changes of digital twin entities. On the other hand, data spaces “powered by FIWARE” can also support the exchange of large files using standard file transfer protocols, since this kind of file transfers may be required for certain scenarios like training of AI algorithms.

Note that systems participating in Data Spaces “powered by FIWARE” do not need to be themselves “powered by FIWARE”. Systems which have not been architected using FIWARE can still use the NGSI API to share data they produce and consume data they need in the form of data associated with attributes of Digital Twin entities which represent that part of the world they deal with. This can be done directly by the systems or through NGSI system adapters which have been programmed to perform a conversion between NGSI and the API that the system natively supports for managing data.

From a theoretical perspective, systems connected to a Data Space should be able to share data using the API they prefer. The specification (information model) of each API could be published as some kind of manifesto that certain components, integrated as part of the platform that

systems should use to connect to the Data Space, can dynamically process in order to perform an automated adaptation from/to the APIs. However, such an approach faces important challenges. In the first place, such an approach has only been demonstrated in very simple scenarios and involving very simple APIs. Thus, the ability to exploit the kind of sophisticated features NGSI-LD would support for accessing Digital Twin data will be rather limited. On the other hand, creating a “common lingua” has proven to work for creating the kind of ecosystems we pursue: we shall ask ourselves if the world wide web had experienced the speed on adoption reached if HTTP and HTML hadn’t been adopted as “common lingua” for web servers and browsers and each web server had the ability to choose a different protocol (as opposed to HTTP) or a different document format (as opposed to HTML). NGSI-LD has the advantages of being an open standard (defined by ETSI), has a strong Open Source community behind (FIWARE) and, quite relevant within Europe, it will pave the way for alignment with developments in the Connecting Europe Facility (CEF) program. Creation of system adapters that transform from specific APIs a given system may still require to use from/to NGSI-LD has proven to be not a complex task.

5 The Smart Data Models initiative

As mentioned before, the NGSI-LD API is domain-agnostic, therefore it is designed to work for any type of Digital Twin. Consequent-

ly, achieving full interoperability requires also the adoption of common data models to be represented in formats compatible with the API.




Figure 8 - Smart Data Models organization on GitHub

Here, the [Smart Data Models initiative](#), already introduced in a previous section, reaches great relevance. It brings a powerful resource for developers who can rely on the way data model specifications are mapped into concrete JSON and JSON-LD structures under the initiative, compatible with NGSIv2 and NGSI-LD respectively.

Figure 8 illustrates how resources are organized within the Smart Data Models initiative on GitHub. Data Models are grouped into “subjects” (weather, parking, aquaculture, etc) which in turn are referred to from repositories associated with the multiple application domains being considered (Smart Cities, Smart AgriFood, Smart Manufacturing, Smart Water, Smart Energy, etc). Note that there are subjects which are very specific to a given application domain (e.g., “street lighting” with regards to smart cities and communities) while oth-

ers may be relevant to multiple domains (e.g., “weather” that is relevant to almost every domain or “sewage” that is relevant to the Smart Cities and Smart Water domains). Published data models are open to contributions and royalty-free.

An Open Governance Model has been defined for the Smart Data Models initiative defining the lifecycle of data models comprising incubation of brand new data models as well as curation of data models via harmonization of different contributions. Processes and procedures for management of the different activities follow best practices from Open Source communities, guided by principles of transparency and meritocracy.

Conclusions

A Digital Twin approach can be the basis for architecting vertical smart solutions, a system of systems approach for digitising smart organizations and the creation of data spaces. FIWARE brings two core aspects that are essential to standardize in the field of Digital Twins, namely the API to be used for management of digital twin data, and the definition of data twin classes. This way, solutions based on the Digital Twin paradigm will be fully interoperable and portable across platforms. FIWARE is bringing today mature open source technologies,

compatible with CEF Building Blocks, that may accelerate the delivery of smart solutions based on the Digital Twin paradigm. The open source nature of FIWARE guarantees that users of such solutions will not be locked-in to any particular vendor. In addition the FIWARE ecosystem has proven the ability to create standards which are accepted and adapted in the meantime on a global scale. Actually, solution providers and end users in the Americas, in Africa, in India or in Asia are building their smart solutions based on FIWARE standards and technologies.

FIWARE FOR **DIGITAL TWINS**

POSITION PAPER
VERSION 0.0 / JUNE 2021



Be certified and featured
in the FIWARE Marketplace.

[GO TO THE MARKETPLACE](#)



Never miss an update.
Join our Newsletter.

[SUBSCRIBE](#)

FIND US ON



Twitter



Facebook



LinkedIn



YouTube



Github